arXiv:1607.02793v2 [math.OC] 21 Mar 2017

On Faster Convergence of Cyclic Block Coordinate Descent-type Methods for Strongly Convex Minimization*

Xingguo Li, Tuo Zhao, Raman Arora, Han Liu, and Mingyi Hong†

Abstract

The cyclic block coordinate descent-type (CBCD-type) methods, which performs iterative updates for a few coordinates (a block) simultaneously throughout the procedure, have shown remarkable computational performance for solving strongly convex minimization problems. Typical applications include many popular statistical machine learning methods such as elastic-net regression, ridge penalized logistic regression, and sparse additive regression. Existing optimization literature has shown that for strongly convex minimization, the CBCD-type methods attain iteration complexity of $\mathcal{O}(p\log(1/\epsilon))$, where ϵ is a pre-specified accuracy of the objective value, and p is the number of blocks. However, such iteration complexity explicitly depends on p, and therefore is at least p times worse than the complexity $\mathcal{O}(\log(1/\epsilon))$ of gradient descent (GD) methods. To bridge this theoretical gap, we propose an improved convergence analysis for the CBCD-type methods. In particular, we first show that for a family of quadratic minimization problems, the iteration complexity $\mathcal{O}(\log^2(p)\cdot\log(1/\epsilon))$ of the CBCD-type methods matches that of the GD methods in term of dependency on p, up to a $\log^2 p$ factor. Thus our complexity bounds are sharper than the existing bounds by at least a factor of $p/\log^2(p)$. We also provide a lower bound to confirm that our improved complexity bounds are tight (up to a $\log^2 p$ factor), under the assumption that the largest and smallest eigenvalues of the Hessian matrix do not scale with p. Finally, we generalize our analysis to other strongly convex minimization problems beyond quadratic ones.

1 Introduction

We consider a class of composite convex minimization problems:

$$x^* = \underset{x\in\mathbb{R}^d}{\operatorname{argmin}}\, \mathcal{F}(x),\ \mathcal{F}(x) = \mathcal{L}(x) + \mathcal{R}(x), \tag{1}$$

*Some preliminary results in this paper were presented at the 19th International Conference on Artificial Intelligence and Statistics (Li et al., 2016).

†Xingguo Li and Tuo Zhao equally contributed to this paper. Manuscript as March 20, 2017. XL is affiliated with Department of Electrical and Computer Engineering at University of Minnesota; TZ is affiliated with School of Industrial and Systems Engineering at Georgia Institute of Technology; RA is affiliated with Department of Computer Science at Johns Hopkins University; HL is affiliated with Department of Operations Research and Financial Engineering, Princeton University; and MH is affiliated with Department of Industrial and Manufacturing Systems Engineering at Iowa State University. Emails: {lixx1661}@umn.edu, {tourzhao}@gatech.edu, {arora}@cs.jhu.edu, {hanliu}@princeton.edu, {mingyi}@iastate.edu.

where $\mathcal{L}(\cdot)$ is a twice differentiable loss function and $\mathcal{R}(\cdot)$ is a possibly nonsmooth and strongly convex penalty function. Many popular statistical machine learning problems are of the form (1), such as elastic-net regression (Zou and Hastie, 2005), ridge penalized logistic regression (Hastie et al., 2009), support vector machine (Vapnik and Vapnik, 1998) and many others (Hastie et al., 2009). For notational simplicity, we assume that there exists a partition of d coordinates such that

$$x = [x_1^\top, \ldots, x_p^\top]^\top \in \mathbb{R}^d,$$

where $x_j \in \mathbb{R}^{d_j}$, $d = \sum_{j=1}^p d_j$, and $d_j \ll p$. The penalty function $\mathcal{R}(x)$ in these applications is block coordinate decomposable, i.e., $\mathcal{R}(x) = \sum_{j=1}^p \mathcal{R}_j(x_j)$. Then we can rewrite the objective in (1) as

$$\mathcal{F}(x) = \mathcal{L}(x_1, \ldots, x_p) + \sum_{j=1}^p \mathcal{R}_j(x_j).$$

Many algorithms such as gradient decent-type (GD-type) methods (Nesterov, 2004, 2007; Beck and Teboulle, 2009a,b; Becker et al., 2011), cyclic block coordinate descent-type (CBCD-type) methods (Luo and Tseng, 1992; Tseng, 1993, 2001; Friedman et al., 2007; Liu et al., 2009; Tseng and Yun, 2009; Saha and Tewari, 2013; Schmidt and Friedlander, 2015; Zhao and Liu, 2015; Zhao et al., 2014b,a, 2012; Li et al., 2015b), and alternating direction method of multipliers (ADMM) (Gabay and Mercier, 1976; Boyd et al., 2011; He and Yuan, 2012; Hong and Luo, 2012; Zhao and Liu, 2012; Liu et al., 2014, 2015; Li et al., 2015a)) have been proposed to solve (1). Among these algorithms, the CBCD-type methods have been immensely successful (Friedman et al., 2007, 2010; Mazumder et al., 2011; Tibshirani et al., 2012; Razaviyayn et al., 2013; Zhao et al., 2014a). One popular instance of the CBCD-type methods is the cyclic block coordinate minimization (CBCM) method, which minimizes (1) with respect to a single block of variables while holding the rest fixed. Particularly, at the $(t+1)$-th iteration, given $x^{(t)}$, we choose to solve a collection of optimization problems: For $j = 1, \ldots, p$,

$$x_j^{(t+1)} = \operatorname*{argmin}_{x_j} \mathcal{L}\Big(x_{1:(j-1)}^{(t+1)}, x_j, x_{(j+1):p}^{(t)}\Big) + \mathcal{R}_j(x_j), \tag{2}$$

where $x_{1:(j-1)}^{(t+1)}$ and $x_{(j+1):p}^{(t)}$ are defined as

$$x_{1:(j-1)}^{(t+1)} = [x_1^{(t+1)\top}, \ldots, x_{j-1}^{(t+1)\top}]^\top \text{ and } x_{(j+1):p}^{(t)} = [x_{j+1}^{(t)\top}, \ldots, x_p^{(t)\top}]^\top.$$

For some applications (e.g. elastic-net penalized linear regression), we can obtain a simple closed form solution to (2), but for many other applications (e.g. ridge-penalized logistic regression), (2) does not admit a closed form solution and requires more sophisticated optimization procedures.

A popular alternative to CBCD-type methods is to solve a quadratic approximation of (2) using the cyclic block coordinate gradient descent (CBCGD) method. For notational simplicity, we denote the partial gradient $\nabla_{x_j}\mathcal{L}(x)$ by $\nabla_j\mathcal{L}(x)$. Then the CBCGD method solves a collection of optimization problems: For $j = 1, \ldots, p$,

$$x_j^{(t+1)} = \operatorname*{argmin}_{x_j} (x_j - x_j^{(t)})^\top \nabla_j \mathcal{L}\Big(x_{1:(j-1)}^{(t+1)}, x_j, x_{(j+1):p}^{(t)}\Big) + \frac{\eta_j}{2}\|x_j - x_j^{(t)}\|^2 + \mathcal{R}_j(x_j), \tag{3}$$

where $\eta_j > 0$ is a step-size parameter for the j-th block.

There have been many results on iteration complexity of block coordinate descent-type (BCD-type) methods, but most of them focus on the randomized BCD-type methods, where blocks are randomly chosen with replacement in each iteration (Shalev-Shwartz and Tewari, 2011; Richtárik and Takáč, 2012; Lu and Xiao, 2015), which demonstrate better iteration complexities than cyclic BCD-type methods in the worst case scenarios (Lee and Wright, 2016; Sun and Ye, 2016). In contrast, existing literature on cyclic BCD-type methods is rather limited. Beck and Tetruashvili (2013) focuses on minimizing smooth objective functions, and has shown that given a pre-specified accuracy ϵ for the objective value, the CBCGD method attains linear iteration complexity of $\mathcal{O}(\log(1/\epsilon))$ for minimizing smooth and strongly convex problems, and sublinear iteration complexity of $\mathcal{O}(1/\epsilon)$ for smooth and nonstrongly convex problems. Hong et al. (2013) focuses on minimizing nonsmooth composite objective functions such as (1), and has shown that the CBCM and CBCGD methods attain sublinear iteration complexity of $\mathcal{O}(1/\epsilon)$, when the objective function is nonstronlgy convex.

Here, we are interested in establishing an improved iteration complexity of the CBCM and CBCGD methods, when the nonsmooth composite objective function is strongly convex. Beck and Tetruashvili (2013) has shown that for smooth minimization, the CBCGD method attains linear iteration complexity of

$$\mathcal{O}\Big(\mu^{-1}L^2 p \log(1/\epsilon)\Big), \tag{4}$$

where L is the Lipschitz constant of the gradient mapping $\nabla\mathcal{L}(x)$ and μ is the strongly convex coefficient of the objective function. However, such iteration complexity depends on p (the number of blocks), and therefore is at least p times worse than the complexity $\mathcal{O}\Big(\mu^{-1}L\log(1/\epsilon)\Big)$ of the gradient descent (GD) methods.

To bridge this theoretical gap, we propose an improved convergence analysis for the CBCD-type methods. Specifically, we show that for a family of quadratic minimization problems, the iteration complexity of the CBCD-type methods matches that of the GD methods in term of dependency on p up to a $\log^2(p)$ factor. More precisely, when $\mathcal{L}(x)$ is quadratic, the iteration complexity of the CBCD-type methods is

$$\mathcal{O}\Big(\mu^{-1}L^2\log^2(p)\log(1/\epsilon)\Big). \tag{5}$$

As can be seen easily, (5) is better than (4) by a factor of $p/\log^2(p)$. We also provide a lower bound analysis that confirms that our improved iteration complexity is tight up to a $\log^2(p)$ factor if the largest and smallest eigenvalues of the Hessian matrix do not scale with p. Finally, we generalize our analysis to other strongly convex minimization problems beyond quadratic minimization. Specifically, for smooth minimization, the iteration complexity of the BCGD method is

$$\mathcal{O}\Big(\mu^{-1}p\log(1/\epsilon)\cdot\min\{L^2,p\}\Big) \tag{6}$$

This indicates that (6) is better than (4) when L^2 is significantly larger than p, e.g. ill-conditioned problems; for more details refer to Table 1. It is worth mentioning that all the above results on the

CBCD-type methods can be used to establish the iteration complexity for the popular permuted BCM (PBCM) and permuted BCGD (PBCGD) methods, in which the blocks are randomly sampled without replacement in each round.

Table 1: Compared with Beck and Tetruashvili (2013), our contributions are manyfold: (1) Developing the iteration complexity bounds of the CBCM and CBCGD methods for different specifications on $\mathcal{L}(\cdot)$ and $\mathcal{R}(\cdot)$; (2) Developing the iteration complexity bound of CBCGD for quadratic $\mathcal{L}(\cdot)$ + nonsmooth $\mathcal{R}(\cdot)$; (3) Improving the iteration complexity bound of CBCGD for smooth $\mathcal{R}(\cdot)$.

	Method	$\mathcal{L}(\cdot)$	$\mathcal{R}(\cdot)$	Our analysis	Beck and Tetruashvili (2013)
[a]	CBCGD	Quadratic	Smooth	$\mathcal{O}\left(\mu^{-1}\log^2(p)L^2\log(1/\epsilon)\right)$	$\mathcal{O}\left(\mu^{-1}pL^2\log(1/\epsilon)\right)$
[b]	CBCGD	Quadratic	Nonsmooth	$\mathcal{O}\left(\mu^{-1}\log^2(p)L^2\log(1/\epsilon)\right)$	N/A
[c]	CBCGD	General Convex	Smooth	$\mathcal{O}\left(\mu^{-1}p\min\{L^2,p\}\log(1/\epsilon)\right)$	$\mathcal{O}\left(\mu^{-1}pL^2\log(1/\epsilon)\right)$
[d]	CBCGD	General Convex	Nonsmooth	$\mathcal{O}\left(\mu^{-1}pL^2\log(1/\epsilon)\right)$	N/A
[e]	CBCM	Quadratic	Smooth	$\mathcal{O}\left(\mu^{-1}\log^2(p)L^2\log(1/\epsilon)\right)$	N/A
[f]	CBCM	Quadratic	Nonsmooth	$\mathcal{O}\left(\mu^{-1}\log^2(p)L^2\log(1/\epsilon)\right)$	N/A
[g]	CBCM	General Convex	Smooth	$\mathcal{O}\left(\mu^{-1}pL^2\log(1/\epsilon)\right)$	N/A
[h]	CBCM	General Convex	Nonsmooth	$\mathcal{O}\left(\mu^{-1}pL^2\log(1/\epsilon)\right)$	N/A

Remark: See Theorem 3 for [a] and [b]; See Theorem 4 for [e] and [f]; See Theorem 7 for [c]; See Theorem 8 for [d], [g], and [h]. When $\mathcal{R}(\cdot)$ is nonsmooth, the optimization problem is actually solved by the cyclic block coordinate proximal gradient (CBCPGD) method. For notational convenience in this paper, however, we simply call it the CBCGD method.

The rest of the paper is organized as follows. In Section 2, we introduce notations and preliminary assumptions. Then we provide the main results of improved convergence analysis for CBCD-type approachs in Section 3. Numerical evaluations are provided in Section 4, followed by further discussions in Section 5.

2 Notations and Assumptions

We start with introducing notations used in this paper. Given a vector $\boldsymbol{v} = (v_1,\ldots,v_d)^\top \in \mathbb{R}^d$, we define vector norms: $\|\boldsymbol{v}\|_1 = \sum_j |v_j|$, $\|\boldsymbol{v}\|^2 = \sum_j v_j^2$, and $\|\boldsymbol{v}\|_\infty = \max_j |v_j|$. Let $\{\mathcal{A}_1,\ldots,\mathcal{A}_p\}$ be a partition of all d coordinates with $|\mathcal{A}_j| = d_j$ and $\sum_{j=1}^p d_j = d$. We use $\boldsymbol{v}_j$ to denote the subvector of $\boldsymbol{v}$ with all indices in $\mathcal{A}_j$. Given a matrix $A \in \mathbb{R}^{d\times d}$, we use $\lambda_{\max}(A)$ and $\lambda_{\min}(A)$ to denote the largest and smallest eigenvalues of A. We denote $\|A\|$ as the spectral norm of A (i.e., the largest singular value). We denote $\otimes$ and $\odot$ as the Kronecker product and Hadamard (entrywise) product for two matrices respectively.

Before we proceed with our analysis, we introduce some assumptions on $\mathcal{L}(\cdot)$ and $\mathcal{R}(\cdot)$.

Assumption 1. *$\mathcal{L}(\cdot)$ is convex, and its gradient mapping $\nabla\mathcal{L}(\cdot)$ is Lipschitz continuous and also blockwise Lipschitz continuous, i.e., there exist positive constants L and L_j's such that for any $x, x' \in \mathbb{R}^d$ and $j = 1, \ldots, p$, we have*

$$\|\nabla\mathcal{L}(x') - \nabla\mathcal{L}(x)\| \leq L\|x - x'\| \;\; \text{and} \;\; \|\nabla_j\mathcal{L}\left(x_{1:(j-1)}, x'_j, x_{(j+1):p}\right) - \nabla_j\mathcal{L}(x)\| \leq L_j\|x_j - x'_j\|.$$

We define $L_{\max} = \max_j L_j$ and $L_{\min} = \min_j L_j$.

Assumption 2. *$\mathcal{R}(\cdot)$ is strongly convex and also blockwise strongly convex, i.e., there exist positive constants μ and μ_j's such that for any $x, x' \in \mathbb{R}^d$ and $j = 1, \ldots, p$, we have*

$$\mathcal{R}(x) \geq \mathcal{R}(x') + (x - x')^\top\xi' + \frac{\mu}{2}\|x - x'\|^2 \;\; \text{and} \;\; \mathcal{R}_j(x_j) \geq \mathcal{R}_j(x'_j) + (x_j - x'_j)^\top\xi'_j + \frac{\mu_j}{2}\|x_j - x'_j\|^2,$$

for all $\xi' \in \partial\mathcal{R}(x')$. We define $\mu_{\min} = \min_j \mu_j$.

For notational simplicity, we define auxiliary variables

$$L^{\mu}_{\min} = \min_j L_j + \mu_j \;\; \text{and} \;\; y^{(t,j)} = [x^{(t)\top}_{1:(j-1)}, x^{(t-1)\top}_{j:p}]^\top, \; j = 1, \ldots, p.$$

Our analysis considers $L_{\min}$, $L_{\max}$, $L^{\mu}_{\min}$, $\mu_{\min}$, μ, and $d_{\max} = \max_j d_j$ as constants, which do not scale with the block size p as in existing literature (Beck and Tetruashvili, 2013).

3 Improved Convergence Analysis

Our analysis consists of the following three steps:

(1) Characterize the successive descent after each CBCD iteration;

(2) Characterize the gap towards the optimal objective value after each CBCD iteration;

(3) Combine (1) and (2) to establish the iteration complexity bound.

We present our analysis under different specifications on $\mathcal{L}(\cdot)$ and $\mathcal{R}(\cdot)$.

3.1 Quadratic Minimization

We first consider a scenario, where $\mathcal{L}(\cdot)$ is a quadratic function. Particularly, we solve

$$x^* = \operatorname*{argmin}_{x\in\mathbb{R}^d} \mathcal{L}(x) + \mathcal{R}(x) = \operatorname*{argmin}_{\substack{x_j\in\mathbb{R}^{d_j} \\ j=1,\ldots,p}} \frac{1}{2}\left\|\sum_{j=1}^{p} A_{*j}x_j - b\right\|^2 + \sum_{j=1}^{p}\mathcal{R}_j(x_j), \tag{7}$$

where $A_{*j} \in \mathbb{R}^{n\times d_j}$ for $j = 1, ..., p$. Typical applications of (7) in statistical machine learning include ridge regression, elastic-net penalized regression, and sparse additive regression.

We first characterize the successive descent of the CBCGD method.

Lemma 1. *Recall that $\mathcal{F}$ is the objective defined in (1). Suppose that Assumptions 1 and 2 hold. We choose $\eta_j = L_j$ for the CBCGD method. Then for all $t \geq 1$, we have*

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) \geq \frac{L^{\mu}_{\min}}{2}\|x^{(t)} - x^{(t+1)}\|^2.$$

Proof. At t-th iteration, there exists a $\xi_j^{(t+1)} \in \partial\mathcal{R}_j(x_j^{(t+1)})$ satisfying the optimality condition for each sub-problem:

$$\nabla_j \mathcal{L}(y^{(t+1,j+1)}) + \eta_j(x_j^{(t+1)} - x_j^{(t)}) + \xi_j^{(t+1)} = 0. \tag{8}$$

Then by definition of CBCGD given in (3), we have

$$\begin{aligned}\mathcal{F}(y^{(t+1,j)}) - \mathcal{F}(y^{(t+1,j+1)}) = (x_j^{(t)} - x_j^{(t+1)})^\top \nabla_j \mathcal{L}(y^{(t+1,j+1)}) + \mathcal{R}_j(x_j^{(t)}) \\ - \frac{L_j}{2}\|x_j^{(t+1)} - x_j^{(t)}\|^2 - \mathcal{R}_j(x_j^{(t+1)}).\end{aligned} \tag{9}$$

By Assumptions 2, we have

$$\mathcal{R}_j(x_j^{(t)}) - \mathcal{R}_j(x_j^{(t+1)}) \geq (x_j^{(t)} - x_j^{(t+1)})^\top \xi_j^{(t+1)} + \frac{\mu_j}{2}\|x_j^{(t)} - x_j^{(t+1)}\|^2. \tag{10}$$

Combining (8), (9), and (10), we have

$$\mathcal{F}(y^{(t+1,j)}) - \mathcal{F}(y^{(t+1,j+1)}) \geq \frac{L_j + \mu_j}{2}\|x_j^{(t)} - x_j^{(t+1)}\|^2. \tag{11}$$

We complete the proof via summation of (11) over $j = 1, \ldots, p$ and the definition of $L^{\mu}_{\min}$. □

Next, we characterize the gap towards the optimal objective value.

Lemma 2. *Suppose that Assumptions 1 and 2 hold. Then for all $t \geq 1$, we have*

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq \frac{L^2 \log^2(2pd_{\max})}{2\mu}\|x^{(t+1)} - x^{(t)}\|^2.$$

Proof. Since $\mathcal{L}(x)$ is quadratic, its second order Taylor expansion is tight, i.e.

$$\mathcal{L}(x^*) = \mathcal{L}(x^{(t+1)}) + \langle \nabla\mathcal{L}(x^{(t+1)}), x^* - x^{(t+1)}\rangle + \frac{1}{2}\|A(x^{(t+1)} - x^*)\|^2, \tag{12}$$

where $A = \left[A_{*1}, \ldots, A_{*p}\right] \in \mathbb{R}^{n\times d}$.

Consider matrices $\widetilde{P}$ and $\widetilde{A}$, defined as

$$\widetilde{P} = \begin{bmatrix} L_1 & 0 & 0 & \ldots & 0 & 0 \\ 0 & L_2 & 0 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & & \ldots & 0 & L_p \end{bmatrix} \in \mathbb{R}^{p\times p} \quad \text{and} \quad \widetilde{A} = \begin{bmatrix} A_{*1} & 0 & 0 & \ldots & 0 & 0 \\ 0 & A_{*2} & 0 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ldots & \vdots & \vdots \\ 0 & 0 & & \ldots & 0 & A_{*p} \end{bmatrix} \in \mathbb{R}^{np\times d}.$$

For simplicity, we assume that $d_1 = ... = d_p = m = d/p$. For any $s \in \mathbb{Z}^+$, we define the lower triangular matrix $D_s \in \mathbb{R}^{s\times s}$ as

$$D_s = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 & 0 \\ 1 & 1 & 0 & \dots & 0 & 0 \\ 1 & 1 & 1 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ 1 & 1 & 1 & \dots & 1 & 1 \end{bmatrix}$$

By the definition of L_j, we have

$$L_j \geq \lambda_{max}(A_j^\top A_j), \forall j.$$

which gives us the following inequality

$$\widetilde{P} \otimes I_n \geq \widetilde{A}^\top \widetilde{A}. \tag{13}$$

To characterize the gap towards the optimal objective, we have

$$\begin{aligned}
&\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \\
&\overset{\text{(i)}}{\leq} \langle \nabla \mathcal{L}(x^{(t+1)}), x^{(t+1)} - x^* \rangle + \mathcal{R}(x^{(t+1)}) - \mathcal{R}(x^*) \\
&\overset{\text{(ii)}}{\leq} \langle \nabla \mathcal{L}(x^{(t+1)}), x^{(t+1)} - x^* \rangle + \langle \xi^{(t+1)}, x^{(t+1)} - x^* \rangle - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2 \\
&\overset{\text{(iii)}}{\leq} \sum_{j=1}^{p} \langle \nabla_j \mathcal{L}(x^{(t+1)}) - \nabla_j \mathcal{L}\left(x^{(t)}\right), x_j^{(t+1)} - x_j^* \rangle - \sum_{j=1}^{p} L_j \langle x_j^{(t+1)} - x_j^{(t)}, x_j^{(t+1)} - x_j^* \rangle - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2 \\
&= \sum_{j=1}^{p} \left\langle \sum_{k \geq j} A_k (x_k^{(t+1)} - x_k^{(t)}), A_j (x_j^{(t+1)} - x_j^*) \right\rangle - (x^{(t+1)} - x^{(t)})^\top (\widetilde{P} \otimes I_m)(x^{(t+1)} - x^*) - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2 \\
&\leq (x^{(t+1)} - x^{(t)})^\top \widetilde{A}^\top (D_p \otimes I_m) \widetilde{A} (x^{(t+1)} - x^*) - (x^{(t+1)} - x^{(t)})^\top (\widetilde{P} \otimes I_m)(x^{(t+1)} - x^*) - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2 \\
&= (x^{(t+1)} - x^{(t)})^\top \left(\widetilde{A}^\top (D_p \otimes I_m) \widetilde{A} - \widetilde{P} \otimes I_m \right)(x^{(t+1)} - x^*) - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2 \\
&\overset{\text{(iv)}}{=} (x^{(t+1)} - x^{(t)})^\top \left(\left(A^\top A - \widetilde{A}^\top \widetilde{A} \right) \odot D_d + \widetilde{A}^\top \widetilde{A} - \widetilde{P} \otimes I_m \right)(x^{(t+1)} - x^*) - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2,
\end{aligned}$$

where (i) is from (12), (ii) is from Assumption 2, (iii) is from the optimality condition to the subproblem associated with x_j,

$$\langle \nabla_j \mathcal{L}(x^{(t)}) + L_j (x_j^{(t+1)} - x_j^{(t)}) + \xi_j^{(t+1)}, x_j - x_j^{(t+1)} \rangle \geq 0 \text{ for any } x_j \in \mathbb{R}^m,$$

and (iv) comes from the fact that

$$\widetilde{A}^\top (D_p \otimes I_m) \widetilde{A} = \left(A^\top A - \widetilde{A}^\top \widetilde{A} \right) \odot D_d + \widetilde{A}^\top \widetilde{A},$$

where $\odot$ denotes the Hadamard product and $\mathbf{1}_n \in \mathbb{R}^{n\times n}$ is a matrix with all entries as 1.

Let us define

$$B = \left(A^\top A - \widetilde{A}^\top \widetilde{A}\right) \odot D_d + \widetilde{A}^\top \widetilde{A} - \widetilde{P} \otimes I_m,$$

then we have

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq (x^{(t+1)} - x^{(t)})^\top B(x^{(t+1)} - x^*) - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2, \tag{14}$$

Minimizing the R.H.S. of the above inequality over x^*, we obtain

$$-\mu(x^* - x^{(t+1)}) - B^\top(x^{(t+1)} - x^{(t)}) = 0.$$

which implies

$$x^* = -\frac{B^\top(x^{(t+1)} - x^{(t)})}{\mu} + x^{(t+1)}. \tag{15}$$

Plugging (15) into (14), we obtain

$$\begin{aligned}
\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) &\leq \frac{1}{2\mu}\|B(x^{(t+1)} - x^{(t)})\|^2 \leq \frac{\lambda_{\max}^2(B)}{2\mu}\|x^{(t+1)} - x^{(t)}\|^2 \\
&\overset{\text{(i)}}{\leq} \frac{\lambda_{\max}^2\left(\left(A^\top A - \widetilde{A}^\top \widetilde{A}\right) \odot D_d\right)}{2\mu}\|x^{(t+1)} - x^{(t)}\|^2 \\
&\overset{\text{(ii)}}{\leq} \frac{\lambda_{\max}^2\left(A^\top A - \widetilde{A}^\top \widetilde{A}\right)\left(1 + \frac{1}{\pi} + \frac{\log(d)}{\pi}\right)^2}{2\mu}\|x^{(t+1)} - x^{(t)}\|^2 \\
&\overset{\text{(iii)}}{\leq} \frac{\lambda_{\max}^2\left(A^\top A - \widetilde{A}^\top \widetilde{A}\right)\log^2(2d)}{2\mu}\|x^{(t+1)} - x^{(t)}\|^2 \\
&\overset{\text{(iv)}}{\leq} \frac{L^2\log^2(2d)}{2\mu}\|x^{(t+1)} - x^{(t)}\|^2,
\end{aligned}$$

where (i) comes from (13), which indicates that $\lambda_{\max}(\widetilde{A}^\top \widetilde{A} - \widetilde{P} \otimes I_m) \leq 0$, (iii) is true if $d \geq 3$, and (iv) comes from $d \leq p \cdot d_{\max}$ and the fact that

$$\lambda_{\max}\left(A^\top A - \widetilde{A}^\top \widetilde{A}\right) \leq \lambda_{\max}\left(A^\top A\right) + \lambda_{\max}\left(-\widetilde{A}^\top \widetilde{A}\right) \leq \lambda_{\max}\left(A^\top A\right) \leq L.$$

Inequality (ii) follows from the result on the spectral norm of the triangular truncation operator in Angelos et al. (1992). More specifically, let us define

$$L_d = \max\left\{\frac{\|A \odot D_d\|}{\|A\|} : A \in \mathbb{R}^{d\times d}, A \neq 0\right\}.$$

Then we have

$$\left|\frac{L_d}{\log d} - \frac{1}{\pi}\right| \leq \frac{\left(1 + \frac{1}{\pi}\right)}{\log d}.$$

The final claim holds by $d \leq p \cdot d_{\max}$.

□

Using Lemmas 1 and 2, we establish the iteration complexity bound of the CBCGD method for minimizing (7) in the next theorem.

Theorem 3. *Suppose that Assumptions 1 and 2 hold. We choose $\eta_j = L_j$ for the CBCGD method. Given a pre-specified accuracy $\epsilon > 0$ of the objective value, we need at most*

$$\left\lceil \frac{\mu L_{\min}^{\mu} + L^2 \log^2(2pd_{\max})}{\mu L_{\min}^{\mu}} \log\left(\frac{\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)}{\epsilon}\right) \right\rceil$$

iterations for the CBCGD method to ensure that $\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \leq \epsilon$

Proof. Combining Lemmas 1 and 2, we obtain

$$\begin{aligned}
\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) &= [\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)})] + [\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*)] \\
&\geq \frac{L_{\min}^{\mu}}{2} \|x^{(t)} - x^{(t+1)}\|^2 + [\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*)] \\
&\geq \left(1 + \frac{L_{\min}^{\mu} \mu}{L^2 \log^2(2pd_{\max})}\right) [\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*)].
\end{aligned}$$

Recursively applying the above inequality for $t \geq 1$, we obtain

$$\frac{\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*)}{\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)} \leq \left(1 - \frac{\mu L_{\min}^{\mu}}{\mu L_{\min}^{\mu} + L^2 \log^2(2pd_{\max})}\right)^t.$$

To ensure $\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \leq \epsilon$, we only need a large enough t to ensure that

$$\left(1 - \frac{\mu L_{\min}^{\mu}}{\mu L_{\min}^{\mu} + L^2 \log^2(2pd_{\max})}\right)^t [\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)] \leq \epsilon.$$

We complete the proof by the above inequality, and the basic inequality $\kappa \geq \log^{-1}\left(\frac{\kappa}{\kappa-1}\right)$. □

As can be seen in Theorem 3, the iteration complexity depends on p only in the order of $\log^2(p)$, which is generally mild in practice. The iteration complexity of the CBCM method can be established in a similar manner.

Theorem 4. *Suppose that Assumptions 1 and 2 hold. Given a pre-specified accuracy ϵ, we need at most*

$$\left\lceil \frac{\mu \mu_{\min} + 4L^2 \log^2(2pd_{\max})}{\mu \mu_{\min}} \log\left(\frac{\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)}{\epsilon}\right) \right\rceil$$

iterations for the CBCM method such that $\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \leq \epsilon$

Proof. The overall proof also consists of three major steps: (i) successive descent, (ii) gap towards the optimal objective value, and (iii) iteration complexity.

Successive Descent: At t-th iteration, there exists a $\xi_j^{(t+1)} \in \partial \mathcal{R}_j(x_j^{(t+1)})$ satisfying the optimality condition:

$$\nabla_j \mathcal{L}(y^{(t+1,j+1)}) + \xi_j^{(t+1)} = 0. \tag{16}$$

Then we have

$$\begin{aligned}\mathcal{F}(y^{(t+1,j)}) - \mathcal{F}(y^{(t+1,j+1)}) &\overset{\text{(i)}}{\geq} (x_j^{(t)} - x_j^{(t+1)})^\top \nabla_j \mathcal{L}(y^{(t+1,j+1)}) + \mathcal{R}_j(x_j^{(t)}) - \mathcal{R}_j(x_j^{(t+1)}) \\ &\overset{\text{(ii)}}{\geq} \left(\nabla_j \mathcal{L}(y^{(t+1,j+1)}) + \xi_j^{(t+1)}\right)^\top \cdot (x_j^{(t)} - x_j^{(t+1)}) + \frac{\mu_j}{2}\|x_j^{(t)} - x_j^{(t+1)}\|^2 \\ &\overset{\text{(iii)}}{=} \frac{\mu_j}{2}\|x_j^{(t)} - x_j^{(t+1)}\|^2, \end{aligned} \tag{17}$$

where (i) is from the convexity of $\mathcal{L}(\cdot)$, (ii) is from Assumptions 2 and $\eta_j \geq L_j$, and (iii) is from (8). By summation of (11) over $j = 1, \ldots, p$, we have

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) \geq \frac{\mu_{\min}}{2}\|x^{(t)} - x^{(t+1)}\|^2.$$

Gap towards the Optimal Objective Value: The proof follows the same arguments with the proof of Lemma 2, with a few differences.

First,with the optimality condition to the subproblem associated with x_j,

$$\langle \nabla_j \mathcal{L}(x^{(t)}) + \xi_j^{(t+1)}, x_j - x_j^{(t+1)} \rangle \geq 0 \text{ for any } x_j \in \mathbb{R}^m,$$

we have

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq (x^{(t+1)} - x^{(t)})^\top B(x^{(t+1)} - x^*) - \frac{\mu}{2}\|x^{(t+1)} - x^*\|^2,$$

where $B = \left(A^\top A - \widetilde{A}^\top \widetilde{A}\right) \odot D_d + \widetilde{A}^\top \widetilde{A}$.

Then, using the same technique to bound the eigenvalues for matrices with Hadamard product, we have

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq \frac{L^2 \log^2(2d) + L_{\max}^2}{\mu}\|x^{(t+1)} - x^{(t)}\|^2 \leq \frac{2L^2 \log^2(2d)}{\mu}\|x^{(t+1)} - x^{(t)}\|^2.$$

Iteration Complexity: The analysis follows from that of Theorem 3. □

Theorem 4 establishes that the iteration complexity of the CBCM method matches that of the CBCGD method. To the best of our knowledge, Theorems 3 and 4 are the sharpest iteration complexity analysis of the CBCD-type methods for minimizing (7).

3.2 The Tightness of the Iteration Complexity for Quadratic Problems

We next provide an example to establish the tightness of the above result. We consider the following optimization problem

$$\min_x \mathcal{H}(x) := \|Bx\|^2, \tag{18}$$

where $B \in \mathbb{R}^{p\times p}$ is a tridiagonal Toplitz matrix defined as follows:

$$B = \begin{bmatrix} 3 & 1 & 0 & 0 & \dots & 0 & 0 & 0 \\ 1 & 3 & 1 & 0 & \dots & 0 & 0 & 0 \\ 0 & 1 & 3 & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 & 3 & 1 \\ 0 & 0 & 0 & 0 & \dots & 0 & 1 & 3 \end{bmatrix}.$$

Note that the minimizer to (18) is $x^* = [0,0,\dots,0]^\top$, and the eigenvalues of B are given by $3 + 2\cos(j\pi/(j+1))$ for $j = 1,\dots,p$. Since the Hessian matrix of (18) is $2B^\top B$, we have

$$L = \lambda_{\max}(2B^\top B) \le 50, \mu = \lambda_{\min}(2B^\top B) \ge 2, \mu_{\min} = 10.$$

Clearly, for this problem the largest and smallest eigenvalues of the Hessian matrix, as well as L/γ do not scale with p. We consider each coordinate $x_j \in \mathbb{R}$ as a block. Then the problem can be rewritten as $\min \|\sum_{j=1}^p B_{*j}x_j\|$, where B_{*j} denotes the j-th column of B. Given an initial solution $x^{(0)}$, we can show that $x^{(1)}$ is generated by

$$x_1^{(1)} = -\frac{1}{4}\left(4x_2^{(0)} + x_3^{(0)}\right), \tag{19}$$

$$x_2^{(1)} = -\frac{1}{5}\left(4x_1^{(1)} + 4x_3^{(0)} + x_4^{(0)}\right) \tag{20}$$

$$x_3^{(1)} = -\frac{1}{5}\left(x_1^{(1)} + 4x_2^{(1)} + x_4^{(0)} + x_5^{(0)}\right), \tag{21}$$

$$x_j^{(1)} = -\frac{1}{5}\left(x_{j-2}^{(1)} + 4x_{j-1}^{(1)} + x_{j+1}^{(0)} + x_{j+2}^{(0)}\right) \tag{22}$$

$$x_{p-1}^{(1)} = -\frac{1}{5}\left(x_{p-3}^{(1)} + 4x_{p-2}^{(0)} + 4x_p^{(0)}\right), \tag{23}$$

$$x_p^{(1)} = -\frac{1}{4}\left(x_{p-2}^{(1)} + 4x_{p-1}^{(1)}\right). \tag{24}$$

Now we choose the initial solution

$$x^{(0)} = \left[1, \frac{9}{32}, \frac{7}{8}, 1, \dots, 1, 1\right]^\top.$$

Then by (19)–(24), we obtain

$$x^{(1)} = \left[-\frac{1}{2}, -\frac{1}{2}, \dots, -\frac{1}{2}, -\frac{3}{10}, -\frac{17}{40}\right]^\top,$$

which yields

$$\mathcal{H}(x^{(1)}) - \mathcal{H}(x^*) \ge \frac{25}{4}(p-3),$$
$$\|x^{(0)} - x^*\|^2 \le p - 2 + \left(\frac{9}{32}\right)^2 + \left(\frac{7}{8}\right)^2 \le p - 1.$$

Therefore, we have

$$\frac{\mathcal{H}(x^{(1)}) - \mathcal{H}(x^*)}{\|x^{(0)} - x^*\|^2} \geq \frac{25(p-3)}{4p} \geq \frac{22}{4}.$$

This implies that when the largest and smallest eigenvalues of the Hessian matrix do not scale with p (the number of blocks), the iteration complexity is independent of p, and cannot be further improved.

3.3 General Smooth Minimization

We next consider general strongly convex smooth minimization, which includes Beck and Tetruashvili (2013) as a special case. Here we require $\mathcal{R}(x)$ to be smooth and strongly convex.

Assumption 3. *$\mathcal{R}(\cdot)$ is smooth and also blockwise smooth, i.e., there exist positive constants β and β_j's such that for $x, x' \in \mathbb{R}^d$ and $j = 1, \ldots, p$, we have*

$$\mathcal{R}(x) \leq \mathcal{R}(x') + (x - x')^\top \nabla \mathcal{R}(x') + \frac{\beta}{2}\|x - x'\|^2 \text{ and}$$
$$\mathcal{R}_j(x_j) \leq \mathcal{R}_j(x_j') + (x_j - x_j')^\top \nabla_j \mathcal{R}(x') + \frac{\beta_j}{2}\|x_j - x_j'\|^2.$$

Moreover, we define $\beta_{\max} = \max_j \beta_j$.

Moreover, we assume that the Hessian matrix H of the objective function $\mathcal{F}$ exists, which is denoted as $H_{ij}(x) = \frac{\partial \mathcal{F}(x)}{\partial x_i \partial x_j}$.

Since the objective function is globally smooth, the CBCGD method can directly take the update form: For $j = 1, \ldots, p$,

$$x_j^{(t+1)} = x_j^{(t)} - \eta_j \left(\nabla_j \mathcal{L}(y^{(t+1,j+1)}) + \nabla \mathcal{R}_j(x_j^{(t)}) \right),$$

where $\eta_j > 0$ is a step-size parameter for the j-th block.

Typical applications of the general strongly convex smooth minimization in statistical machine learning includes ridge penalized logistic regression, and ridge penalized multinomial regression. It is worth mentioning that our analysis for the general case is applicable to smooth quadratic minimization, but is very different from the analysis in previous sections for quadratic minimization. For simplicity, we only consider $d_{\max} = d_1 = \ldots = d_p = 1$, and leave the generalization to $d_{\max} > 1$ to future work.

We first characterize the successive descent after each coordinate gradient descent (CGD) iteration.

Lemma 5. *Suppose that Assumptions 1 and 3 hold. We choose $\eta_j = L_j + \beta_j$ for the CBCGD method. Then for all $t \geq 1$, there exists $z^{(t,j)}$ in the line segment of $(x^{(t)}, y^{(t,j)})$ for each $j \in \{1, \ldots, p\}$ such that*

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) \geq \frac{\|\nabla \mathcal{F}(x^{(t)})\|^2}{2\left(L_{\max}^{\beta} + \frac{\|H\|^2}{L_{\min}^{\beta}}\right)},$$

where H is defined as

$$H \triangleq \begin{bmatrix} 0 & 0 & 0 & \dots & 0 & 0 \\ H_{21} & 0 & 0 & \dots & 0 & 0 \\ H_{31} & H_{32} & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \dots & \ddots & \vdots \\ H_{p1} & H_{p2} & H_{p3} & \dots & H_{p,p-1} & 0 \end{bmatrix}, \tag{25}$$

with $H_{j,i} \triangleq H_{j,i}(\boldsymbol{z}^{(t,j)}) = \frac{\partial \mathcal{F}(\boldsymbol{z}^{(t,j)})}{\partial z_j \partial z_i}$, and $L_{\min}^{\beta}$ and $L_{\max}^{\beta}$ are defined as

$$L_{\min}^{\beta} = \min\{L_j + \mu_j, j = 1, \dots, p\},$$
$$L_{\max}^{\beta} = \max\{L_j + \mu_j, j = 1, \dots, p\}.$$

Proof. We first provide a lower bound of the successive descent using the gradient of $\mathcal{F}(\cdot)$ based on the Lipschitz continuity of $\nabla\mathcal{F}(\cdot)$. We have that $y^{(t,j)}$ and $y^{(t,k+1)}$ only differ at the k-th coordinate, and $\nabla_j \mathcal{F}(y^{t,k})$ has Lipschitz gradient with Lipschitz constant F_j, which implied

$$\begin{aligned} \mathcal{F}(y^{(t,j+1)}) &\le \mathcal{F}(y^{(t,j)}) + (y^{(t,j+1)} - y^{+})^{\top} \nabla_j \mathcal{F}(y^{(t,j)}) + \frac{F_j}{2} \|y^{(t,j+1)} - y^{(t,j)}\|^2 \\ &\overset{\text{(i)}}{=} \mathcal{F}(y^{(t,j)}) - \frac{2L_j^{\beta} - F_j}{2(L_j^{\beta})^2} \|\nabla_j \mathcal{F}(y^{(t,j)})\|^2 \overset{\text{(ii)}}{\le} \mathcal{F}(y^{(t,j)}) - \frac{1}{2L_j^{\beta}} \|\nabla_j \mathcal{F}(y^{(t,j)})\|^2, \end{aligned}$$

where (i) is from that $x_j^{(t+1)} = x_j^{(t)} - \frac{\nabla_j \mathcal{F}(y^{(t,j)})}{L_j^{\beta}}$, and (ii) is from the fact that $L_j^{\beta} \ge F_j$. Then the decrease of the objective is

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) = \sum_{k=1}^{p} \mathcal{F}(y^{(t,j)}) - \mathcal{F}(y^{(t,j+1)}) \ge \sum_{k=1}^{p} \frac{1}{2L_j^{\beta}} \|\nabla_j \mathcal{F}(y^{(t,j)})\|^2. \tag{26}$$

There exists $\boldsymbol{z}^{(t,j)}$ such that

$$\begin{aligned} \nabla_j \mathcal{F}(x^{(t)}) &= \nabla_j \mathcal{F}(x^{(t)}) - \nabla_j \mathcal{F}(y^{(t,j)}) + \nabla_j \mathcal{F}(y^{(t,j)}) \\ &\overset{\text{(i)}}{=} \nabla(\nabla_j \mathcal{F}(\boldsymbol{z}^{(t,j)}))^{\top}(x^{(t)} - y^{(t,j)}) + \nabla_j \mathcal{F}(y^{(t,j)}) \\ &= \left[\frac{\partial \mathcal{F}(\boldsymbol{z}^{(t,j)})}{\partial z_j \partial z_1}, \dots, \frac{\partial \mathcal{F}(\boldsymbol{z}^{(t,j)})}{\partial z_j \partial z_{j-1}}, 0, \dots, 0\right] \left[\frac{x_1^{(t)} - x_1^{(t+1)}}{L_1^{\beta}}, \dots, \frac{x_{j-1}^{(t)} - x_{j-1}^{(t+1)}}{L_{j-1}^{\beta}}, 0, \dots, 0\right]^{\top} + \nabla_j \mathcal{F}(y^{(t,j)}) \\ &= \left[\frac{H_{j1}}{\sqrt{L_1^{\beta}}}, \dots, \frac{H_{j,j-1}}{\sqrt{L_{j-1}^{\beta}}}, 0, \dots, 0\right] \left[\frac{x_1^{(t)} - x_1^{(t+1)}}{\sqrt{L_1^{\beta}}}, \dots, \frac{x_{j-1}^{(t)} - x_{j-1}^{(t+1)}}{\sqrt{L_{j-1}^{\beta}}}, 0, \dots, 0\right]^{\top} + \nabla_j \mathcal{F}(y^{(t,j)}) \\ &= \left[\frac{H_{j1}}{\sqrt{L_1^{\beta}}}, \dots, \frac{H_{j,j-1}}{\sqrt{L_{j-1}^{\beta}}}, \sqrt{L_j^{\beta}}, 0, \dots, 0\right] \left[\frac{\nabla_1 \mathcal{F}(y^{(t,1)})}{\sqrt{L_1^{\beta}}}, \dots, \frac{\nabla_p \mathcal{F}(y^{(t,p)})}{\sqrt{L_p^{\beta}}}\right]^{\top} \\ &= h_j^{\top} f \end{aligned}$$

where (i) is from the mean-value theorem, $h_j = \left[\frac{H_{j1}}{\sqrt{L_1^\beta}}, \ldots, \frac{H_{j,j-1}}{\sqrt{L_{j-1}^\beta}}, \sqrt{L_j^\beta}, 0, \ldots, 0\right]$ and $f = \left[\frac{\nabla_1 \mathcal{F}(y^{(t,1)})}{\sqrt{L_1^\beta}}, \ldots, \frac{\nabla_p \mathcal{F}(y^{(t,p)})}{\sqrt{L_p^\beta}}\right]$. Let $\widetilde{H}$ be

$$\widetilde{H} = \begin{bmatrix} h_1^\top \\ \vdots \\ h_p^\top \end{bmatrix} = \begin{bmatrix} \sqrt{L_1^\beta} & 0 & 0 & \ldots & 0 & 0 \\ \frac{H_{21}(z^{(t,2)})}{\sqrt{L_1^\beta}} & \sqrt{L_2^\beta} & 0 & \ldots & 0 & 0 \\ \frac{H_{31}(z^{(t,3)})}{\sqrt{L_1^\beta}} & \frac{H_{32}(z^{(t,3)})}{\sqrt{L_2^\beta}} & 0 & \ldots & 0 & 0 \\ \vdots & \vdots & \vdots & \ldots & \ddots & \vdots \\ \frac{H_{p1}(z^{(t,p)})}{\sqrt{L_1^\beta}} & \frac{H_{p2}(z^{(t,p)})}{\sqrt{L_2^\beta}} & \frac{H_{p3}(z^{(t,p)})}{\sqrt{L_3^\beta}} & \ldots & \frac{H_{p,p-1}(z^{(t,p)})}{\sqrt{L_{p-1}^\beta}} & \sqrt{L_p^\beta} \end{bmatrix}$$

Then we have

$$\begin{aligned} \|\nabla \mathcal{F}(x^{(t)})\|^2 &= \sum_{j=1}^p \|\nabla_j \mathcal{F}(x^{(t)})\|^2 = \sum_{j=1}^p \|h_j^\top f\|^2 = \|\widetilde{H} f\|^2 \\ &\le \|\widetilde{H}\|^2 \|f\|^2 = \|\widetilde{H}\|^2 \sum_{k=1}^p \frac{1}{2L_j^\beta} \|\nabla_j \mathcal{F}(y^{(t,j)})\|^2. \end{aligned} \tag{27}$$

Let $\widetilde{P}$ be defined as in the proof of Lemma 2. Then we have

$$\|\widetilde{H}\|^2 = \|\widetilde{P}^{1/2} + H\widetilde{P}^{-1/2}\|^2 \le 2\left(\|\widetilde{P}^{1/2}\|^2 + \|H\widetilde{P}^{-1/2}\|^2\right) \le 2\left(L_{\max}^\beta + \frac{\|H\|^2}{L_{\min}^\beta}\right), \tag{28}$$

Combining (26), (27) and (28), we have

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) \ge \sum_{k=1}^p \frac{1}{2L_j^\beta} \|\nabla_j \mathcal{F}(y^{(t,j)})\|^2 \ge \frac{\|\nabla \mathcal{F}(x^{(t)})\|^2}{\|\widetilde{H}\|^2} \ge \frac{\|\nabla \mathcal{F}(x^{(t)})\|^2}{2\left(L_{\max}^\beta + \frac{\|H\|^2}{L_{\min}^\beta}\right)}.$$

□

We now characterize the gap towards the optimal objective after each CGD iteration.

Lemma 6. *Suppose that Assumptions 1 and 2 hold. Then, for all $t \ge 1$, we have*

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \le \frac{\|\nabla \mathcal{F}(x^{(t)})\|^2}{2\mu}.$$

Proof. From the convexity of $\mathcal{L}(\cdot)$ and strong convexity of $\mathcal{R}(\cdot)$, we have

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \le (x^{(t)} - x^*)^\top \nabla \mathcal{F}(x^{(t)}) - \frac{\mu}{2}\|x^{(t)} - x^*\|^2.$$

Minimizing the right hand side over x^*, we have $x^* = x^{(t)} - \frac{\nabla \mathcal{F}(x^{(t)})}{\mu}$ and the desired result. □

Combining the two lemmas above, we establish the iteration complexity bound of CGD.

Theorem 7. *Suppose that Assumption 1, 2 and 3 hold. We choose $\eta_j = L_j + \beta_j$. Then, given a pre-specified accuracy ϵ, we need at most*

$$\left\lceil \frac{\mu + L_{\max} + \frac{\min\{pL^2, p^2L_{\max}^2\}}{L_{\min}}}{\mu} \log\left(\frac{\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)}{\epsilon}\right) \right\rceil$$

iterations such that $\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \leq \epsilon$.

Proof. We first bound $\|H\|$. Since L_i is the Lipschitz constant of $\nabla_i \mathcal{L}$, we have $|H_{ji}| \leq \sqrt{L_j L_i}$, which further implies

$$\|H\|^2 \leq \|H\|_F^2 = \sum_{j=2}^{p} \sum_{i=1}^{j-1} |H_{ji}|^2 \leq \sum_{j=2}^{p} \sum_{i=1}^{j-1} L_j L_i \leq p^2 L_{\max}^2.$$

An alternative bound is

$$\|H\|^2 \leq \|H\|_F^2 = \sum_{j=1}^{p} \sum_{i=1}^{p} H_{ji}^2 \leq \sum_{j=1}^{p} L^2 \leq pL^2.$$

Thus we have $\|H\| \leq \min\{p^2 L_{\max}^2, pL^2\}$. We only need to combine Lemmas 5 and 6, and complete the proof by following similar lines to the proof of Theorem 3. □

As can be seen from Theorem 7, the established iteration complexity bound is sharper than that in Beck and Tetruashvili (2013) for ill-conditioned optimization problems, where we often observe $L^2 \gg p$.

3.4 General Nonsmooth Minimization

We provide an iteration complexity bound of the CBCM and CBCGD methods for a general $\mathcal{L}(\cdot)$ and a nonsmooth $\mathcal{R}(\cdot)$.

Theorem 8. *Suppose that Assumptions 1 and 2 hold. We choose $\eta_j = L_j$ for the CBCGD method. Then given a pre-specified accuracy ϵ of the objective value, we need at most*

$$\left\lceil \frac{\mu L_{\min}^{\mu} + 4pL^2}{\mu L_{\min}^{\mu}} \log\left(\frac{\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)}{\epsilon}\right) \right\rceil$$

iterations for the CBCGD method and at most

$$\left\lceil \frac{\mu\mu_{\min} + pL^2}{\mu\mu_{\min}} \log\left(\frac{\mathcal{F}(x^{(0)}) - \mathcal{F}(x^*)}{\epsilon}\right) \right\rceil$$

iterations for the CBCM method to guarantee $\mathcal{F}(x^{(t)}) - \mathcal{F}(x^*) \leq \epsilon$.

Proof. **Successive Descent**: For CBCGD, using the same analysis of Lemma 1, we have that for all $t \geq 1$,

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) \geq \frac{L^{\mu}_{\min}}{2}\|x^{(t)} - x^{(t+1)}\|^2.$$

For CBCM, using the same analysis of Theorem 4, we have that for all $t \geq 1$,

$$\mathcal{F}(x^{(t)}) - \mathcal{F}(x^{(t+1)}) \geq \frac{\mu_{\min}}{2}\|x^{(t)} - x^{(t+1)}\|^2.$$

Gap towards the Optimal Objective Value: By the strong convexity of $\mathcal{R}(\cdot)$, we have

$$\mathcal{F}(x) - \mathcal{F}(x^{(t+1)}) \geq \frac{\mu}{2}\|x - x^{(t+1)}\|^2 + (x - x^{(t+1)})^{\top}(\nabla\mathcal{L}(x^{(t+1)}) + \xi^{(t+1)}), \tag{29}$$

where $\xi_j^{(t+1)} \in \partial\mathcal{R}_j(x_j^{(t+1)})$. We then minimize both sides of (29) with respect to x and obtain

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq \frac{\|\nabla\mathcal{L}(x^{(t+1)}) + \xi^{(t+1)}\|^2}{2\mu}, \tag{30}$$

For CBCGD, we have

$$\begin{aligned}
\|\nabla\mathcal{L}(x^{(t+1)}) + \xi^{(t+1)}\|^2 &\overset{\text{(i)}}{\leq} \sum_{j=1}^{p}\|\nabla\mathcal{L}(x^{(t+1)}) - \nabla_j\mathcal{L}(y^{(t+1,j+1)}) - L_j(x_j^{(t+1)} - x_j^{(t)})\|^2 \\
&\leq \sum_{j=1}^{p} 2\|\nabla\mathcal{L}(x^{(t+1)}) - \nabla_j\mathcal{L}(y^{(t+1,j+1)})\|^2 + 2L_j^2\|x_j^{(t+1)} - y_j^{(t+1,j)}\|^2 \\
&\leq 4pL^2\|x^{(t+1)} - x^{(t)}\|^2,
\end{aligned} \tag{31}$$

where (i) comes from the optimality condition

$$\nabla_j\mathcal{L}(y^{(t+1,j+1)}) + L_j(x_j^{(t+1)} - x_j^{(t)}) + \xi_j^{(t+1)} = 0.$$

Combining (30) and (31), we have

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq \frac{2pL^2\|x^{(t+1)} - x^{(t)}\|^2}{\mu}.$$

For CBCM, we have

$$\|\nabla\mathcal{L}(x^{(t+1)}) + \xi^{(t+1)}\|^2 \overset{\text{(i)}}{\leq} \sum_{j=1}^{p}\|\nabla\mathcal{L}(x^{(t+1)}) - \nabla_j\mathcal{L}(y^{(t+1,j+1)})\|^2 \leq pL^2\|x^{(t+1)} - x^{(t)}\|^2, \tag{32}$$

where (i) comes from the optimality condition

$$\nabla_j\mathcal{L}(y^{(t+1,j+1)}) + \xi_j^{(t+1)} = 0.$$

Combining (30) and (32), we have

$$\mathcal{F}(x^{(t+1)}) - \mathcal{F}(x^*) \leq \frac{pL^2\|x^{(t+1)} - x^{(t)}\|^2}{2\mu}.$$

Iteration Complexity: The analysis follows from that of Theorem 3. □

Theorem 8 is a general result for general nonsmooth minimization. In contrast, Beck and Tetruashvili (2013) only covers general smooth minimization.

3.5 Extensions to Nonstrongly Convex Minimization

For nonstrongly convex minimization, we only need to add a strongly convex perturbation to the objective function

$$\widehat{x} = \operatorname{argmin} \mathcal{F}(x) + \frac{\sigma}{2}\|x\|^2, \tag{33}$$

where $\sigma > 0$ is a perturbation parameter. Then, the results above can be used to analyze the CBCD-type methods for minimizing (33). Eventually, by setting σ as a reasonable small value, we can establish $\mathcal{O}(1/\epsilon)$-type iteration complexity bounds up to a $\log(1/\epsilon)$ factor. See Shalev-Shwartz and Zhang (2014) for more details.

4 Numerical Results

We consider two typical statistical machine learning problems as examples to illustrate our analysis.

(I) Elastic-net Penalized Linear Regression: Let $A \in \mathbb{R}^{n\times d}$ be the design matrix, and $b \in \mathbb{R}^n$ be the response vector. We solve the following optimization problem

$$\min_{x\in\mathbb{R}^d} \ \frac{1}{2n}\|b - Ax\|^2 + \lambda_1\|x\|^2 + \lambda_2\|x\|_1,$$

where λ is the regularization parameter. We set $n = 10{,}000$ and $d = 20{,}000$. We simply treat each coordinate as a block (i.e., $d_{\max} = 1$). Each row of A is independently sampled from a 20,000-dimensional Gaussian distribution with mean 0 and covariance matrix Σ. We randomly select 2,000 entries of x, each of which is independently sampled from a uniform distribution over support $(-2, +2)$. The response vector b is generated by the linear model $b = Ax+\epsilon$, where ϵ is sampled from an n-variate Gaussian distribution $N(0, I_n)$. We set $\lambda_1 = \sqrt{\log 1/n}$ and $\lambda_2 = \sqrt{\log d/n} \approx 0.0315$. We normalize A to have $\|A_{*j}\| = \sqrt{n}$ for $j = 1,..,d$, where A_{*j} denotes the j-th column of A. For the BCGD method, we choose $\eta_j = 1$. For the gradient descent method, we either choose $\eta = \lambda_{\max}\left(\frac{1}{n}A^\top A\right)$, or adaptively select η by backtracking line search.

(II) Ridge Penalized Logistic Regression: We solve the following optimization problem

$$\min_{x\in\mathbb{R}^d} \ \frac{1}{n}\sum_{i=1}^{n}\Big[\log(1+\exp(x^\top A_{i*})) - b_i x^\top A_{i*}\Big] + \lambda\|x\|^2.$$

We generate the design matrix A and regression coefficient vector x using the same scheme as sparse linear regression. Again we treat each coordinate as a block (i.e., $d_{\max} = 1$). The response $b = [b_1,...,b_n]^\top$ is generated by the logistic model $b_i = \text{Bernoulli}([1+\exp(-x^\top A_{*i})]^{-1})$. We set $\lambda = \sqrt{1/n}$. For the BCGD method, we choose $\eta_j = \frac{1}{4}$. For gradient descent methods, we choose either the step size $\eta = \frac{1}{4}\lambda_{\max}\left(\frac{1}{n}A^\top A\right)$ or adaptively select η by backtracking line search.

We evaluate the computational performance using the number of passes over p blocks of coordinates (normalized iteration complexity). For the CBCGD method, we count one iteration as one pass (all p blocks). For the randomized BCGD (RBCGD) method, we count p iterations as one

pass (since it only updates one block in each iteration). Besides the CBCGD and RBCGD methods, we also consider a variant of the CBCGD method named the permuted BCGD (PBCGD) method, which randomly permutes all indices for the p blocks in each iteration. Since the RBCGD and PBCGD methods are inherently stochastic, we report the objective values averaged over 20 different runs. Moreover, for the RBCGD method, the block of coordinates is selected uniformly at random in each iteration. We consider four different settings for both elastic-net penalized linear regression and ridge penalized logistic regression based on different choices of the covariance matrix Σ for generating the design matrix. We always choose $\Sigma_{jj} = 1$ for $j = 1, \ldots, d$, and for any $k \neq j$, we set (I) $\Sigma_{jk} = 0$; (II) $\Sigma_{jk} = 0.5$; (III) $\Sigma_{jk} = 0.75$; (IV) $\Sigma_{jk} = 0.5^{|j-k|}$. Note that the condition number of the Hessian matrix depends on Σ. Setting (I) and (IV) tend to yield well-conditioned Hessian matrices whereas Settings (II) and (III) tend to yield a badly conditioned Hessian matrix.

Figure 1 plots the gap between the objective value and the optimal as a function of number of passes for different methods. Our empirical findings can be summarized as follows: (1) All BCD-type methods attain better performance than the GD methods; (2) When the Hessian matrix is ill conditioned, i.e., in Setting (II) and (III), the CBCGD performs worse than the RBCGD and PBCGD methods, which suggests that there is a gap between cyclic and randomized BCGD. (3) When the Hessian matrix is well conditioned (e.g., in Settings (I) and (IV)), all three BCD-type methods attain good performance, and the CBCGD method slightly outperforms the PBCGD method; (4) The CBCGD method outperforms the RBCGD method in Setting (IV).

5 Discussions

Existing literature has established an iteration complexity of $\mathcal{O}(\mu^{-1}L \cdot \log(1/\epsilon))$ for the gradient descent methods when solving strongly convex composite problems. However, our analysis shows that the CBCD-type methods only attains an iteration complexity of $\mathcal{O}(\mu^{-1}pL^2 \cdot \log(1/\epsilon))$. Even though our analysis further shows that the iteration complexity of the CBCD-type methods can be further improved to $\mathcal{O}(\mu^{-1}\log^2(p)L^2 \cdot \log(1/\epsilon))$ for a quadratic $\mathcal{L}(\cdot)$, there still exists a gap of factor $L\log^2 p$. As our numerical experiments show, however, the CBCD-type methods can actually attain a better computational performance than the gradient methods regardless of whether $\mathcal{L}(\cdot)$ is quadratic or not, thereby suggesting that perhaps there is still room for improvement in the iteration complexity analysis of the CBCD-type methods.

It is also worth mentioning that though some literature claims that the CBCD-type methods works as well as the randomized BCD-type methods in practice, there do exist some counter examples, e.g. our experiment in Setting (I), where the CBCD-type methods fail significantly. This suggests that the CBCD-type methods do have some possible disadvantages in practice. To the best of our knowledge, we are not aware of any similar experimental results reported in existing literature.

Furthermore, our numerical results show that the permuted BCD-type methods, which can be viewed as a hybrid of the cyclic and the randomized BCD-type (RBCD-type) methods, has a stable performance irrespective of the problem being well conditioned or not. But to the best of our knowledge, no iteration complexity result has been established for the permuted BCD-type

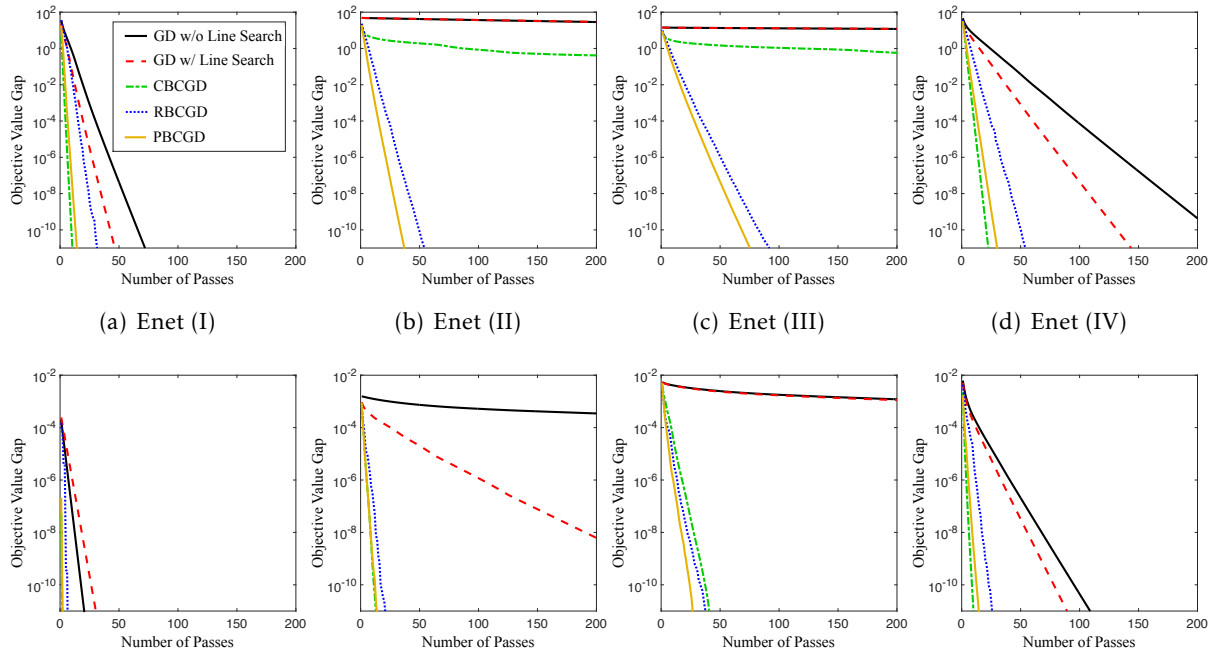

(a) Enet (I) (b) Enet (II) (c) Enet (III) (d) Enet (IV)

(e) Logit (I) (f) Logit (II) (g) Logit (III) (h) Logit (IV)

Figure 1: Comparison among different methods under different settings. "RBCGD" and "PBCGD" denote the randomized BCD-type and permuted BCD-type methods respectively. The vertical axis corresponds to the gap towards the optimal objective value, $\log[\mathcal{F}(x) - \mathcal{F}(x^*)]$; the horizontal axis corresponds to the number of passes over p blocks of coordinates. Though all methods attain linear iteration complexity, their empirical behaviors are different from each others. Note that in plot (b) the curves for the CBCGD method and the RBCGD methods overlap.

(PBCD-type) methods. We leave these problems for future investigation.

References

Angelos, J. R., Cowen, C. C. and Narayan, S. K. (1992). Triangular truncation and finding the norm of a hadamard multiplier. *Linear algebra and its applications* **170** 117–135.

Beck, A. and Teboulle, M. (2009a). Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems. *Image Processing, IEEE Transactions on* **18** 2419–2434.

Beck, A. and Teboulle, M. (2009b). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences* **2** 183–202.

Beck, A. and Tetruashvili, L. (2013). On the convergence of block coordinate descent type methods. *SIAM Journal on Optimization* **23** 2037–2060.

Becker, S., Bobin, J. and Candès, E. J. (2011). Nesta: A fast and accurate first-order method for sparse recovery. *SIAM Journal on Imaging Sciences* **4** 1–39.

Boyd, S., Parikh, N., Chu, E., Peleato, B. and Eckstein, J. (2011). Distributed optimization and statistical

learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning* **3** 1–122.

Friedman, J., Hastie, T., Höfling, H. and Tibshirani, R. (2007). Pathwise coordinate optimization. *The Annals of Applied Statistics* **1** 302–332.

Friedman, J., Hastie, T. and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software* **33** 1–13.

Gabay, D. and Mercier, B. (1976). A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications* **2** 17–40.

Hastie, T., Tibshirani, R., Friedman, J., Hastie, T., Friedman, J. and Tibshirani, R. (2009). *The elements of statistical learning*. Springer.

He, B. and Yuan, X. (2012). On non-ergodic convergence rate of douglas-rachford alternating direction method of multipliers. Tech. rep., Nanjing University.

Hong, M. and Luo, Z.-Q. (2012). On the linear convergence of the alternating direction method of multipliers. *Mathematical Programming* 1–35.

Hong, M., Wang, X., Razaviyayn, M. and Luo, Z.-Q. (2013). Iteration complexity analysis of block coordinate descent methods. *arXiv preprint arXiv:1310.6957* .

Lee, C.-P. and Wright, S. J. (2016). Random permutations fix a worst case for cyclic coordinate descent. *arXiv preprint arXiv:1607.08320* .

Li, X., Zhao, T., Arora, R., Liu, H. and Hong, M. (2016). An improved convergence analysis of cyclic block coordinate descent-type methods for strongly convex minimization. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*.

Li, X., Zhao, T., Yuan, X. and Liu, H. (2015a). The flare package for high dimensional linear regression and precision matrix estimation in R. *The Journal of Machine Learning Research* **16** 553–557.

Li, X., Zhao, T., Zhang, T. and Liu, H. (2015b). The picasso package for nonconvex regularized m-estimation in high dimensions in R. *Technical Report* .

Liu, H., Palatucci, M. and Zhang, J. (2009). Blockwise coordinate descent procedures for the multi-task lasso, with applications to neural semantic basis discovery. In *Proceedings of the 26th Annual International Conference on Machine Learning*.

Liu, H., Wang, L. and Zhao, T. (2014). Sparse covariance matrix estimation with eigenvalue constraints. *Journal of Computational and Graphical Statistics* **23** 439–459.

Liu, H., Wang, L. and Zhao, T. (2015). Calibrated multivariate regression with application to neural semantic basis discovery. *Journal of Machine Learning Research* **16** 1579–1606.

Lu, Z. and Xiao, L. (2015). On the complexity analysis of randomized block-coordinate descent methods. *Mathematical Programming* **152** 615–642.

Luo, Z.-Q. and Tseng, P. (1992). On the linear convergence of descent methods for convex essentially smooth minimization. *SIAM Journal on Control and Optimization* **30** 408–425.

Mazumder, R., Friedman, J. H. and Hastie, T. (2011). Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association* **106** 1125–1138.

Nesterov, Y. (2004). *Introductory lectures on convex optimization: A basic course*, vol. 87. Springer.

Nesterov, Y. (2007). Gradient methods for minimizing composite objective function. Tech. rep., Université catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Razaviyayn, M., Hong, M. and Luo, Z.-Q. (2013). A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization* **23** 1126–1153.

Richtárik, P. and Takáč, M. (2012). Iteration complexity of randomized block-coordinate descent methods for minimizing a composite function. *Mathematical Programming* 1–38.

Saha, A. and Tewari, A. (2013). On the nonasymptotic convergence of cyclic coordinate descent methods. *SIAM Journal on Optimization* **23** 576–601.

Schmidt, M. and Friedlander, M. (2015). Coordinate descent converges faster with the gauss-southwell rule than random selection. *arXiv preprint arXiv:1506.00552* .

Shalev-Shwartz, S. and Tewari, A. (2011). Stochastic methods for ℓ_1-regularized loss minimization. *The Journal of Machine Learning Research* **12** 1865–1892.

Shalev-Shwartz, S. and Zhang, T. (2014). Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming* 1–41.

Sun, R. and Ye, Y. (2016). Worst-case complexity of cyclic coordinate descent: $o(n^2)$ gap with randomized version. *arXiv preprint arXiv:1604.07130* .

Tibshirani, R., Bien, J., Friedman, J., Hastie, T., Simon, N., Taylor, J. and Tibshirani, R. J. (2012). Strong rules for discarding predictors in lasso-type problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **74** 245–266.

Tseng, P. (1993). Dual coordinate ascent methods for non-strictly convex minimization. *Mathematical programming* **59** 231–247.

Tseng, P. (2001). Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of optimization theory and applications* **109** 475–494.

Tseng, P. and Yun, S. (2009). A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming* **117** 387–423.

Vapnik, V. N. and Vapnik, V. (1998). *Statistical learning theory*, vol. 1. Wiley New York.

Zhao, T. and Liu, H. (2012). Sparse additive machine. In *International Conference on Artificial Intelligence and Statistics*.

Zhao, T. and Liu, H. (2015). Accelerated path-following iterative shrinkage thresholding algorithm with application to semiparametric graph estimation.

Zhao, T., Liu, H., Roeder, K., Lafferty, J. and Wasserman, L. (2012). The huge package for high-dimensional undirected graph estimation in R. *The Journal of Machine Learning Research* **13** 1059–1062.

Zhao, T., Liu, H. and Zhang, T. (2014a). A general theory of pathwise coordinate optimization. *arXiv preprint arXiv:1412.7477* .

Zhao, T., Yu, M., Wang, Y., Arora, R. and Liu, H. (2014b). Accelerated mini-batch randomized block coordinate descent method. In *Advances in neural information processing systems*.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67** 301–320.